

# Learning Arbitrary Pairwise Potentials in CRFs for Semantic Segmentation

Måns Larsson<sup>1</sup> Anurag Arnab<sup>3</sup> Fredrik Kahl<sup>1,2</sup> Shuai Zheng<sup>3</sup> Philip Torr<sup>3</sup>  
<sup>1</sup>Chalmers Univ. of Technology <sup>2</sup>Lund Univ. <sup>3</sup>Oxford Univ.

## Abstract

Are we using the right potential functions in the Conditional Random Field models that are popular in the Vision community? Semantic segmentation and other pixel-level labelling tasks have made significant progress recently due to the deep learning paradigm. However, most state-of-the-art structured prediction methods also include a random field model with a hand-crafted Gaussian potential to model spatial priors, label consistencies and feature-based image conditioning.

In this paper, we challenge this view by developing a new inference and learning framework which can learn arbitrary pairwise CRF potentials. Both standard spatial and high-dimensional bilateral kernels are considered. Our framework is based on the observation that CRF inference can be achieved via projected gradient descent and consequently, can easily be integrated in deep neural networks to allow for end-to-end training. It is empirically demonstrated that such learned potentials can improve segmentation accuracy and that certain label class interactions are indeed better modelled by a non-Gaussian potential. Our framework is evaluated on several public benchmarks for semantic segmentation with improved performance compared to previous state-of-the-art CNN+CRF models.

## 1. Introduction

Markov Random Fields (MRFs), Conditional Random Fields (CRFs) and more generally, probabilistic graphical models are a ubiquitous tool used in a variety of domains spanning Computer Vision, Computer Graphics and Image Processing [24, 4]. In this paper, we focus on the application of MRFs for Computer Vision problems involving per-pixel labelling such as image segmentation. There are many successful approaches in this line of research, such as the interactive segmentation of [32] using graph cuts and the semantic segmentation works of [26, 36] where the parallel mean-field inference algorithm was applied for fast inference. Recently, Convolutional Neural Networks (CNNs) have dominated the field in a variety of recognition tasks [19, 34, 31]. However, we observe that leading segmentation approaches still include CRFs, either as a post-




Figure 1. Learned spatially invariant CRF filters for CITYSCAPES. These filters model contextual relationships between classes and their values can be understood as the energy added when setting one pixel to the first class (e.g., vegetation) and the other pixel with relative position (x-shift, y-shift) to the second class (e.g., traffic sign). Note how the terrain-traffic light filter favours vertical edges.

processing step [10, 11, 17, 9], or as part of the deep neural network itself [40, 27, 2, 28, 23].

We leverage on the idea of embedding inference of graphical models into a neural network. An early example of this idea was presented in [7] where the authors back propagated through the Viterbi algorithm when designing a document recognition system. Similar to [40, 2, 3, 38], we use a recurrent neural network to unroll the iterative inference steps of a CRF. This was first used in [40] and [33] to imitate mean-field inference and to train a fully convolutional network [29, 10] along with a CRF end-to-end via back propagation. In contrast to mean-field, we do not optimize the KL-divergence. Instead, we use a gradient descent approach for the inference that directly minimises the Gibbs energy of the random field and hence avoids the approximations of mean-field. A similar framework was recently suggested in [3] for multi-label classification problems in machine learning with impressive results. However, [3] uses a Structured SVM approach for training whereas we do back propagation through the actual steps of the gradient descent method. Moreover, [15] have recently shown that one can obtain lower energies compared to mean-field inference using gradient descent based optimization schemes. Still, we lack formal algorithmic guarantees of the solution quality compared to, e.g., graph cuts [8].

In many works, the pairwise potentials consist of

parametrized Gaussians [25, 40, 2] and it is only the parameters of this Gaussian which are learned. Our framework can learn arbitrary pairwise potentials which need not be Gaussian, cf. Fig. 1. An early work which learned potentials of a linear chain CRF for sequence modelling is [30]. In [12], a general framework for learning arbitrary potentials in deep structured models was proposed based on approximate Maximum Likelihood learning. One of the advantages with that framework is that data likelihood is maximized in the learning process. However, this involves approximating the partition function which is otherwise intractable. This hinders the handling of large structured output spaces like in our case. The experiments of [12] and [3] are limited to multi-label classification where there is no spatial relationship between different labels as in pixel-labelling tasks. Moreover, [23] have also observed that the memory requirements of the method of [12] renders it infeasible for the large datasets common in Computer Vision. Another approach to learning arbitrary pairwise potentials was presented in [23] which uses Gibbs sampling. Again they struggle with the difficulty of computing the partition function. In the end, only experiments on synthetic data restricted to learned 2D potentials are presented.

The authors of [27] and [9] also learn arbitrary pairwise potentials to model contextual relations between parts of the image. However, their approaches still perform post-processing with a CRF model with parametric Gaussian potentials. In [21], a pairwise potential is learned based on sparse bilateral filtering. Applying such a filter can be regarded as one iteration in the CRF inference step. In [21], the bilateral filter is applied twice, mimicking the first two iterations. Our method is not restricted to a limited number of iterations. Perhaps more importantly is that we not only learn sparse high-dimensional bilateral filters, but also learn arbitrary spatial filters. Such spatial 2D potentials are computationally much more efficient and easier to analyze and interpret compared to their high-dimensional counterparts. We also note that [16] proposed back propagating through mean-field inference to learn parameters. However, this was not in the context of neural networks as in the aforementioned approaches and our work. For pixel-labelling tasks, we focus on discrete random fields. We note that learning arbitrary pairwise potentials for deep structured models with continuous valued output variables has recently been explored by [38].

In summary, our contributions are as follows.

- Our main contribution is a new framework for non-parametric CRF inference and learning which is integrated with standard CNNs. To our knowledge, we are the first to simultaneously learn arbitrary, non-sparse spatial and sparse bilateral filter kernels end-to-end. During inference, we directly minimize the CRF energy using gradient descent and during training, we

back propagate through the gradient descent steps for end-to-end learning.

- We analyze the learned filter kernels empirically and demonstrate that in many cases it is advantageous with non-Gaussian potentials.
- We experimentally compare our approach to several leading methodologies, e.g., [40, 17, 29] and improve on state of the art on two public benchmarks: NYU V2 [37] and CITYSCAPES [14].

Our framework has been implemented in both Caffe [22] and MATCONVNET [35], and all source code will be made publicly available to facilitate further research.

## 2. CRF Formulation

Consider a Conditional Random Field over  $N$  discrete random variables  $\mathcal{X} = \{X_1, \dots, X_N\}$  conditioned on an observation  $\mathbf{I}$  and let  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$  be an undirected graph whose vertices are the random variables  $\{X_1, \dots, X_N\}$ . Each random variable corresponds to a pixel in the image and takes values from a predefined set of  $L$  labels  $\mathcal{L} = \{0, \dots, L-1\}$ . The pair  $(\mathcal{X}, \mathbf{I})$  is modelled as a CRF characterized by the Gibbs distribution

$$P(\mathcal{X} = \mathbf{x} | \mathbf{I}) = \frac{1}{Z(\mathbf{I})} \exp(-E(\mathbf{x} | \mathbf{I})), \quad (1)$$

where  $E(\mathbf{x} | \mathbf{I})$  denotes the Gibbs energy function with respect to the labeling  $\mathbf{x} \in \mathcal{L}^N$  and  $Z(\mathbf{I})$  is the partition function. To simplify notation the conditioning on  $\mathbf{I}$  will from now on be dropped. The MAP inference problem for the CRF model is equivalent to the problem of minimizing the energy  $E(\mathbf{x})$ . In this paper, we only consider energies containing unary and pairwise terms. The energy function can hence be written as

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \psi_i(x_i) + \sum_{(i,j) \in \mathcal{E}} \psi_{ij}(x_i, x_j) \quad (2)$$

where  $\psi_i : \mathcal{L} \rightarrow \mathbb{R}$  and  $\psi_{ij} : \mathcal{L} \times \mathcal{L} \rightarrow \mathbb{R}$  are the unary and pairwise potentials, respectively.

### 2.1. Potentials

The unary potential  $\psi_i(x_i)$  specifies the energy cost of assigning label  $x_i$  to pixel  $i$ . In this work we obtain our unary potentials from a CNN. Roughly speaking, the CNN outputs a probability estimate of each pixel containing each class. Denoting the output of the CNN for pixel  $i$  and class  $x_i$  as  $z_{i:x_i}$ , the unary potential is

$$\psi_i(x_i) = -w_u \log(z_{i:x_i} + \epsilon) \quad (3)$$

where  $w_u$  is a parameter controlling the impact of the unary potentials, and  $\epsilon$  is introduced to avoid numerical problems.

The pairwise potential  $\psi_{ij}(x_i, x_j)$  specifies the energy cost of assigning label  $x_i$  to pixel  $i$  while pixel  $j$  is assigned label  $x_j$ . Introducing pairwise terms in our model enables us to take dependencies between output data into account. We consider the following set of pairwise potentials

$$\psi_{ij}(x_i, x_j) = k_{x_i, x_j}^{spatial}(\mathbf{p}_i - \mathbf{p}_j) + k_{x_i, x_j}^{bilateral}(\mathbf{f}_i - \mathbf{f}_j) \quad (4)$$

Here  $k_{x_i, x_j}^{spatial}$  denotes a spatial kernel with compact support. Its value depends on the relative position coordinates  $\mathbf{p}_i - \mathbf{p}_j$  between pixels  $i$  and  $j$ . We do not restrict these spatial terms to any specific shape, cf. Fig. 1. However, we restrict the support of the potential meaning that if pixels  $i$  and  $j$  are far apart, then the value of  $k_{x_i, x_j}^{spatial}(\mathbf{p}_i - \mathbf{p}_j)$  will be zero. CRFs with Gaussian potentials do not in theory have compact support, and therefore, they are often referred to as dense. However, in practice, the exponential function in the kernel drops off quickly and effectively, the interactions between pixels far apart are negligible.

The term  $k_{x_i, x_j}^{bilateral}$  is a bilateral kernel which depends on the feature vectors  $\mathbf{f}_i$  and  $\mathbf{f}_j$  for pixels  $i$  and  $j$ , respectively. Following several previous works, we let the vector depend on pixel coordinates  $\mathbf{p}_i$  and RGB values associated to the pixel, hence  $\mathbf{f}_i$  is a 5-dimensional vector. Note that for both the spatial and the bilateral kernels, there is one kernel for each label-to-label ( $x_i$  and  $x_j$ ) interaction to enable the model to learn differently shaped kernels for each of these interactions.

## 2.2. Multi-label Graph Expansion and Relaxation

To facilitate a continuous relaxation of the energy minimisation problem we start off by expanding our original graph in the following manner. Each vertex in the original graph  $\mathcal{G}$  will now be represented by  $L$  vertices  $X_{i:\lambda}$ ,  $\lambda \in \mathcal{L}$ . In this way, an assignment of labels in  $\mathcal{L}$  to each variable  $X_i$  is equivalent to an assignment of boolean labels 0 or 1 to each node  $X_{i:\lambda}$ , whereby an assignment of label 1 to  $X_{i:\lambda}$  means that in the multi-label assignment,  $X_i$  receives label  $\lambda$ . To ensure that only one label is assigned to each node, an additional constraint is needed saying that, for each  $i$ , only one of  $X_{i:\lambda}$  are allowed to be labeled 1. This enables to rewrite the energy minimization problem  $\min E(\mathbf{x})$  as the following equivalent integer program

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{V}, \lambda \in \mathcal{L}} \psi_i(\lambda) x_{i:\lambda} + \sum_{\substack{(i,j) \in \mathcal{E} \\ \lambda, \mu \in \mathcal{L}}} \psi_{ij}(\lambda, \mu) x_{i:\lambda} x_{j:\mu} \\ \text{s.t.} \quad & x_{i:\lambda} \in \{0, 1\} \quad \forall i \in \mathcal{V}, \lambda \in \mathcal{L} \\ & \sum_{\lambda \in \mathcal{L}} x_{i:\lambda} = 1 \quad \forall i \in \mathcal{V}. \end{aligned} \quad (5)$$

As a next step, we relax the integer program by allowing real values on the unit interval  $[0, 1]$  instead of booleans only. We denote the relaxed variables  $q_{i:\lambda} \in [0, 1]$ . We can now write our problem as a quadratic program

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{V}, \lambda \in \mathcal{L}} \psi_i(\lambda) q_{i:\lambda} + \sum_{\substack{(i,j) \in \mathcal{E} \\ \lambda, \mu \in \mathcal{L}}} \psi_{ij}(\lambda, \mu) q_{i:\lambda} q_{j:\mu} \\ \text{s.t.} \quad & q_{i:\lambda} \geq 0 \quad \forall i \in \mathcal{V}, \lambda \in \mathcal{L} \\ & \sum_{\lambda \in \mathcal{L}} q_{i:\lambda} = 1 \quad \forall i \in \mathcal{V}. \end{aligned} \quad (6)$$

A natural question is what happens when the domain is enlarged. Somewhat surprisingly, the relaxation is tight [6].

**Proposition 2.1.** *Let  $E(\mathbf{x}^*)$  and  $E(\mathbf{q}^*)$  denote the optimal values of (5) and (6), respectively. Then,*

$$E(\mathbf{x}^*) = E(\mathbf{q}^*).$$

In the supplementary material, we show that for *any* real  $\mathbf{q}$ , one can obtain a binary  $\mathbf{x}$  such that  $E(\mathbf{x}) \leq E(\mathbf{q})$ . In particular, it will be true for  $\mathbf{x}^*$  and  $\mathbf{q}^*$ , which implies  $E(\mathbf{x}^*) = E(\mathbf{q}^*)$ . Note that the proof is constructive.

In summary, it has been shown that to minimize the energy function  $E(\mathbf{x})$  over  $\mathbf{x} \in \mathcal{L}^N$ , one may work in the continuous domain, minimize over  $\mathbf{q}$ , and then replace any solution  $\mathbf{q}$  by a discrete solution  $\mathbf{x}$  which has lower or equal energy. It will only be possible to find a local solution  $\mathbf{q}$ , but still the discrete solution  $\mathbf{x}$  will be no worse than  $\mathbf{q}$ .

## 3. Minimization with Gradient Descent

To solve the program stated in (6) we propose an optimization scheme based on projected gradient descent, see Algorithm 1. It was designed with an extra condition in mind, that all operations should be differentiable to enable back propagation during training.

**Initialize**  $\mathbf{q}^0$   
**for**  $t$  from 0 to  $T - 1$  **do**  
    Compute the gradient  $\nabla_{\mathbf{q}} E(\mathbf{q}^t)$ .  
    Take a step in the negative direction,  
     $\tilde{\mathbf{q}}^{t+1} = \mathbf{q}^t - \gamma \nabla_{\mathbf{q}} E$ .  
    Project  $\tilde{\mathbf{q}}^{t+1}$  to the simplex  $\Delta^L$  satisfying  
     $\sum_{\lambda \in \mathcal{L}} \tilde{q}_{i:\lambda} = 1$  and  $0 \leq \tilde{q}_{i:\lambda} \leq 1$ ,  
     $\mathbf{q}^{t+1} = \text{Proj}_{\Delta^L}(\tilde{\mathbf{q}})$ .  
**end**  
**Output:**  $\mathbf{q}^{T-1}$

Algorithm 1. Projected gradient descent algorithm.

The gradient  $\nabla_{\mathbf{q}} E$  of the objective function  $E(\mathbf{q})$  in (6) has the following elements

$$\frac{\partial E}{\partial q_{i:\lambda}} = \psi_i(\lambda) + \sum_{\substack{j:(i,j) \in \mathcal{E} \\ \mu \in \mathcal{L}}} \psi_{ij}(\lambda, \mu) q_{j:\mu}. \quad (7)$$

The contribution from the spatial kernel in  $\psi_{ij}$ , cf. (4), can be written as

$$v_{i:\lambda}^{spatial} = \sum_{\substack{j:(i,j) \in \mathcal{E} \\ \mu \in \mathcal{L}}} k_{\lambda,\mu}^{spatial} (\mathbf{p}_i - \mathbf{p}_j) q_{j:\mu}. \quad (8)$$

Since the value of the kernel  $v_{i:\lambda}^{spatial}$  only depends on the relative position of pixels  $i$  and  $j$ , the contribution for all pixels and classes can be calculated by passing  $q_{j:\mu}$  through a standard convolution layer consisting of  $L \times L$  filters of size  $(2s + 1) \times (2s + 1)$  where  $L$  is the number of labels and  $s$  the number of neighbours each pixel interacts with in each dimension.

The contribution from the bilateral term is

$$v_{i:\lambda}^{bilateral} = \sum_{\substack{j:(i,j) \in \mathcal{E} \\ \mu \in \mathcal{L}}} k_{\lambda,\mu}^{bilateral} (\mathbf{f}_i - \mathbf{f}_j) q_{j:\mu}. \quad (9)$$

For this computation we utilize the method presented by Jampani *et al.* [21] which is based on the permutohedral lattice introduced by Adams *et al.* [1]. Efficient computations are obtained by using the fact that the feature space is generally sparsely populated. Similar to the spatial filter we get  $L \times L$  filters, each having size of  $(s + 1)^{d+1} - s^{d+1}$  where  $s$  is the number of neighbours each pixel interacts with in each dimension in the sparse feature space.

Next, we simply take a step in the negative direction of the gradient according to

$$\tilde{\mathbf{q}}^{t+1} = \mathbf{q}^t - \gamma \nabla_{\mathbf{q}} E, \quad (10)$$

where  $\gamma$  is the step size.

Finally, we want to project our values onto the simplex  $\Delta^L$  satisfying  $\sum_{\lambda \in \mathcal{L}} q_{i:\lambda} = 1$  and  $0 \leq q_{i:\lambda} \leq 1$ . This is done following the method by Chen *et al.* [13]. Note that this projection is done individually for each pixel  $i$ .

**Comparison to Mean-Field.** In recent years, a popular choice for CRF inference is to apply the mean-field algorithm. One reason is that the kernel evaluations can be computed with fast bilateral filtering [26]. As we have seen in this section, it can be accomplished with our framework as well, with formulas that are less involved. The main difference is that our framework directly optimizes the Gibbs energy which corresponds to MAP while mean-field optimizes KL-divergence which does not.

## 4. Integration in a Deep Learning Framework

In this section we will describe how the gradient descent steps of Algorithm 1 can be formulated as layers in a neural network. We need to be able to calculate error derivatives with respect to the input given error derivatives with respect to the output. In addition we need to be able to calculate

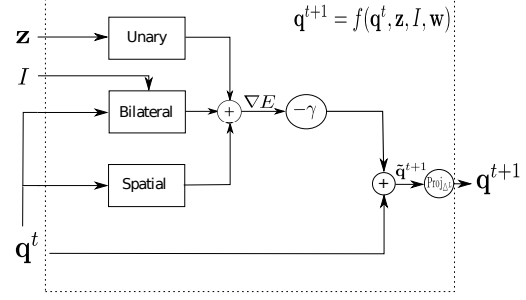


Figure 2. The data flow of one iteration of the projected gradient descent algorithm. Each rectangle or circle represent an operation that can be performed within a deep learning framework.

the error derivatives with respect to the network parameters. This will enable us to unroll the entire gradient descent process as a Recurrent Neural Network (RNN). A schematic of the data flow for one step is shown in Fig. 2. In the supplementary material, all derivative formulas are given.

**Initialization.** The variables  $\mathbf{q}^0$  are set as the output of the CNN, which has been pre-trained to estimate the probability of each pixel containing each class and has a softmax layers as the last layer to ensure that the variables lies within zero and one.

**Gradient Computations.** We have previously explained the gradient computations in Section 3 for the forward pass. To describe the calculation of the error derivatives we first notice that the gradient is calculated by summing three terms, the unary, spatial and bilateral pairwise terms. We can hence treat these three terms separately and combine them using element-wise summing.

The unary term in (3) is an element-wise operation with the CNN output as input and the unary weight  $w_u$  as parameter. The operation is obviously differentiable with respect to both the layer input as well as its parameter. Note that for  $w_u$  we get a summation over all class and pixel indexes for the error derivatives while for the input the error derivatives are calculated element-wise. The spatial pairwise term of the gradient can be calculated efficiently using standard 2D convolution. In addition to giving us an efficient way of performing the forward pass we can also utilize the 2D convolution layer to perform the backward pass, calculating the error derivatives with respect to the input and parameters. Similar to the spatial term, the bilateral term is also calculated utilizing a bilateral filtering technique. Jampani *et al.* [21] also presented a way to calculate the error derivatives with respect to the parameters for an arbitrary shaped bilateral filter.






Figure 3. The data flow of the deep structured model. Each rectangle or circle represent an operation that can be performed within a deep learning framework.

**Gradient Step.** Taking a step in the negative direction of the projected gradient is easily incorporated in a deep learning framework by using an element-wise summing layer. The layer takes the variables  $q^t$  as the first input and the projected gradient (scaled by  $-\gamma$ ) as the second input.

**Simplex Projection** As a final step, the variables from the gradient step  $\tilde{q}^t$  are projected onto the simplex  $\triangle^L$ . In reality we use a leaky version of the last step of the projection algorithm to avoid error derivatives becoming zero during back propagation. Since the projection is done individually for each pixel it can be described as a function  $f(\tilde{q}) : \mathbb{R}^L \rightarrow \mathbb{R}^L$  of which we can calculate the Jacobian, see supplementary materials. Knowing the Jacobian, the error derivatives with respect to the input can be computed during back propagation.

## 5. Recurrent Formulation as a Deep Structured Model

Our iterative solution to the CRF energy minimisation problem by projected gradient descent, as described in the previous sections, is formulated as a Recurrent Neural Network (RNN). The input to the RNN is the image, and the outputs of a CNN, as shown in Fig. 3. The CNN’s output,  $z$ , are the unary potentials and obtained after the final softmax layer (since the CNN is initially trained for classification).

Each iteration of the RNN performs one projected gradient descent step to approximately solve (6). Thus, one update step can be represented by:

$$q^{t+1} = f(q^t, z, I, w). \quad (11)$$

As illustrated in Fig. 3, the gating function  $G_1$  sets  $q^t$  to  $z$  at the first time step, and to  $q^{t-1}$  at all other time steps. In our iterative energy minimisation formulated as an RNN, the output of one step is the input to the next step. We initialise at  $t = 0$  with the output of the unary CNN.

The output of the RNN can be read off  $q^T$  where  $T$  is the total number of steps taken. In practice, we perform a set number of  $T$  steps where  $T$  is a hyperparameter. It is possible to run the RNN until convergence for each image (thus a variable number of iterations per image), but we observed minimal benefit in the final Intersection over Union (IoU) from doing so, as opposed to  $T = 5$  iterations.

The parameters of the RNN are the filter weights for the spatial and bilateral kernels, and also the weight for the unary terms. Since we are able to compute error derivatives with respect to the parameters, and input of the RNN, we can back propagate error derivatives through our RNN to the preceding CNN and train our entire network end-to-end. Furthermore, since the operations of the RNN are formulated as filtering, training and inference can be performed in a fully-convolutional manner.

The CNN part of our network allows us to leverage the ability of CNNs to learn rich feature representations from data, whilst the RNN part of the network utilises the CRF’s ability to model output structure. As we learn the parameters of our pairwise terms, we are not restricted to Gaussian potentials as in [25, 40], and we show the benefits of this in our experiments (Section 6).

**Implementation Details.** Our proposed CRF model has been implemented in the Caffe [22] library, and also has a Matlab wrapper allowing it to be used in MATCONVNET [35]. The Unary CNN part of our model is initialised from a pre-trained segmentation network. As we show in our experiments, different pre-trained networks for different applications can be used.

The CRF model has several tunable parameters. The step size  $\gamma$  and the number of iterations  $T$  specify the properties of the gradient descent algorithm. Too high a step size  $\gamma$  might make the algorithm not end up in a minimum while setting a low step size and a low number of iterations might not give the algorithm a chance to converge. The kernel sizes for the spatial and bilateral kernels also need to be set. Choosing the value of these parameters gives a trade-off between model expression ability and number of parameters, which may cause (or hinder) over-fitting.

The spatial weights of the CRF model are all initialized as zero with the motivation that we did not want to impose a shape for these filters, but instead see what was learned during training. The bilateral filters were initialized as Gaussians with the common Potts class interaction (the filters corresponding to interactions between the same class were set to zero) as done in [26, 10, 40]. Note that unlike [26, 10] we are not limited to only Potts class interactions.

## 6. Experiments

We evaluate the proposed approach on three datasets: WEIZMANN HORSE dataset [5], NYU V2 geometric dataset [37] and CITYSCAPES [14]. In these experiments, we show that the proposed approach, denoted CRF-Grad, has advantages over baseline approaches such as CRFas-RNN [40] and complement other networks such as FCN-8s [29] and LRR [17].




Figure 4. Learned spatially invariant CRF filters of different sizes for the WEIZMANN HORSE dataset. The filters are for the interaction background - horse. The CRF-Grad layer and FCN-8s have been jointly trained in an end-to-end manner.

Method	Mean IoU (%)
FCN-8s (only)	80.0
FCN-8s + spatial (Gaussian)	81.3
FCN-8s + spatial	82.0
FCN-8s + full (Gaussian)	82.9
FCN-8s + full	<b>84.0</b>

Table 1. Quantitative results on the WEIZMANN HORSE dataset. In the entry denoted “full” the complete CRF-Grad layer was used, while in “spatial” no bilateral kernel was used.

### 6.1. Weizmann Horse

The WEIZMANN HORSE dataset is widely used for benchmarking object segmentation algorithms. It contains 328 images of horses in different environments. We divide these images into a training set of 150 images, a validation set of 50 images and a test set of 128 images. Our purpose is to verify our ability to learn reasonable kernels and study the effects of different settings on a relatively small dataset.

The CNN part of our model was initialized as an FCN-8s network [29] pre-trained without the CRF layer. We then compare several variants of our model. We start off by training a variant of our CRF model only using the 2D spatial kernel. We compare these results to using a Gaussian spatial filter, where the parameters for the Gaussian kernel were evaluated using cross-validation. In addition we train the full model with both the spatial and bilateral kernels, once keeping the filters fixed as Gaussians and once allowing arbitrarily shaped filters. All models were trained end-to-end with learning rate of  $10^{-3}$  (normalized by the number of pixels), momentum 0.9 and weight decay  $5 \cdot 10^{-3}$ . The size of the spatial filters for these runs were  $9 \times 9$ .

Quantitative results, mean intersection over union for the test set, are shown in Table 1. Our results show that allowing arbitrarily shaped filters gives better performance than keeping the filters fixed as Gaussian. Also, adding bilateral filters improve the results compared to using spatial filters only.

Fig. 4 shows a visualization of the spatial filters our model learns end-to-end for different kernel sizes. These filters are for the spatial kernels corresponding to the label interaction of class “background” and “horse”. An element

Method	Mean IoU (%)
R-CNN [18]	40.3
Semantic HCRF [37]	42.7
Joint HCRF [37]	44.2
FCN-8s [29]	49.3
Modular CNN [20]	54.3
CRFasRNN [40]	54.4
CRF-Grad bilateral (Ours)	<b>55.0</b>

Table 2. Quantitative comparison between our method and other alternatives on the semantic image segmentation task of NYU V2.

of  $k^{spatial}$  can be understood as the energy added when setting one pixel to “background” and the pixel with the relative position (x-shift,y-shift) to “horse”. The trend we noticed was that for large kernels, the filters learned were similar to Gaussians. However, for smaller kernels, the learned filters had different shapes in many cases.

Fig. 6 presents a convergence analysis for the full version of our layer. The results show that the CRF energy converges in only a few iteration steps and that increasing the number of iterations barely affects performance after  $T = 6$  iterations. The intermediate states of the layer ( $q^t$  for each gradient descent step) are shown in Fig. 5.

### 6.2. NYU V2

The NYU V2 dataset contains images taken by Microsoft Kinect V-1 camera in 464 indoor scenes. We use the official training and validation splits consisting of 795 and 654 images, respectively. Following the setting described in Wang *et al.* [37], we also include additional images for training. These are the images from the NYU V1 dataset that do not overlap with the images in the official validation set. This gives a total of 894 images with semantic label annotations for training. As in [37] we consider 5 classes conveying strong geometric properties: ground, vertical, ceiling, furniture and objects. The CNN part of our model was initialized as the fully convolutional network FCN-8s [29] pre-trained on the data. Afterwards we added our CRF-Grad layer and trained the model end-to-end. The model was trained with learning rate  $1 \cdot 10^{-11}$  (not normalized), momentum 0.99, weight decay 0.005 and batch size 10. The size of the spatial filters were set to  $9 \times 9$  and the bilateral filters to  $s = 1$ . The number of iterations were set to 5, step size to 0.5 and unary weight was initialized as 0.5.

As shown in Table 2, we achieved superior results for semantic image segmentation on the NYU V2 dataset. Some example segmentations are shown in Fig. 7. Additional examples are included in the supplementary material.

### 6.3. Cityscapes

The CITYSCAPES dataset [14] consists of a set of images of street scenes collected from 50 different cities. The




Figure 5. Visualization of intermediate states of the CRF-Grad layer for the WEIZMANN HORSE dataset. Note how each step of the gradient descent algorithm refines the segmentation slightly, removing spurious outlier pixels classified as horse.  $z_{\text{horse}}$  is the CNN output.




Figure 6. Left: CRF energy, as defined in (6), as a function of iterations of the gradient descent method for the WEIZMANN HORSE test set. The thin blue lines show the different instances while the thicker red line show the mean. For these calculations the leak factor was set to zero, meaning that the solutions satisfy the constraints of (6). Note that, for presentation purposes, all energies have been normalized to have the same final energy. Right: Mean IoU as a function of iterations for the WEIZMANN HORSE test data. Note that the performance does not increase after six iterations. The number of iterations were set to five during training.

images are high resolution ( $1024 \times 2048$ ) and are paired with pixel-level annotations of 19 classes including road, sidewalk, traffic sign, pole, building, vegetation and sky. The training, validation and test sets consist of 2975, 500 and 1525 images, respectively. In addition there are 20000 coarsely annotated images that can be used for training.

The CNN part of our model was initialized as an LRR network [17] pre-trained on both the fine and the coarse annotations. We then added our CRF-Grad layer and trained the model end-to-end on the finely annotated images only. The model was trained with a learning rate of  $10^{-3}$  (normalized by the number of pixels), momentum 0.9 and weight decay  $5 \cdot 10^{-3}$ . The size of the spatial filters for these runs were  $9 \times 9$  and the bilateral filters to  $s = 1$ . The number of iterations were set to 5, step size to 0.5 and unary weight was initialized as 0.5.

In Table 3 the results of evaluating our model on the test

Method	Mean IoU (%)
CRFasRNN <sup>*</sup> [40]	62.5
Dilation10 [39]	67.1
Deeplabv2-CRF [11]	70.4
Adelaide_context [27]	71.6
LRR-4x [17]	71.8
CRF-Grad (Ours)	<b>71.9</b>

<sup>\*</sup>Note that CRFasRNN uses a different CNN model than ours.

Table 3. Comparison of our approach and current state-of-the-art methods on the CITYSCAPES test set. Note that only published submissions are included.

set are compared to current state of the art. As can be seen, our model is on par although the improvement upon LRR is minor. An interesting aspect of CITYSCAPES is that it contains classes of thin and vertical objects, e.g., traffic light and pole. What we noticed is that the spatial filters for these classes usually get a more oblong shape. This type of pairwise filters does not add as much energy for switching classes going in the horizontal direction, favoring vertically elongated segmentations. This can be seen in the spatial filter for the class interaction between “terrain” and “traffic light” in Fig. 1. Some example segmentations are shown in Fig. 8. Additional examples as well as class-wise results are included in the supplementary material. Our method got better, or equal, results for 14 of the 19 classes compared to the baseline.

## 7. Concluding Remarks

In this paper, we have introduced a new framework capable of learning arbitrary potentials in random fields models. In a number of experiments, we have empirically demonstrated that our developed framework can improve state-of-the-art CNNs by adding a CRF layer. We have also seen that the learned filters are not necessarily Gaussian, and may

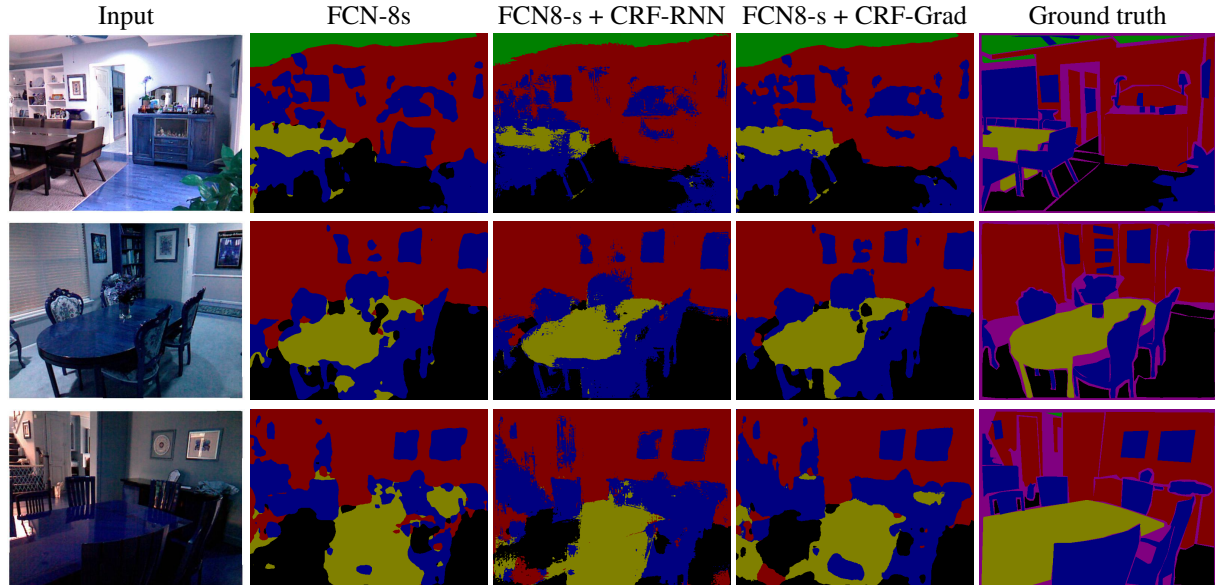


Figure 7. Qualitative results on the NYU V2 dataset. Note that the CRF-Grad captures the shape of the object instances better compared to the baselines. This effect is perhaps most pronounced for the paintings hanging on the walls.




Figure 8. Qualitative results on the CITYSCAPES validation set. Black regions in the ground truth are ignored during evaluation. Our CRF models contextual relationships between classes, hence unlike LRR, it does not label “road” as being on top of “sidewalk” (Row 2). Note that the traffic lights are better segmented with the additional CRF-Grad layer. Adding the CRF-Grad layer increased the IoU of the class traffic lights from 66.8 to 68.1.

capture other kinds of interactions between labels.

A key factor for the success of deep learning and by now a well-established paradigm is that the power of convolutions should be used, especially for the first layers in a CNN. Our work supports that repeated usage of convolutions in the final layers is also beneficial. We also note that our gradient descent steps resemble the highly successful RESNET [19], as one step in gradient descent is, in principle, an identity transformation plus a correction term.

There are several future research avenues that we intend

to explore. In our model, many free variables are introduced and this may lead to over-fitting. One way to compensate would be to collect larger datasets and consider data augmentation. An alternative approach would be to directly encode geometric shape priors into the random fields and thereby reducing the required amount of data.



## References

- [1] A. Adams, J. Baek, and M. A. Davis. Fast high-dimensional filtering using the permutohedral lattice. *Computer Graphics Forum*, 2010. 4
- [2] A. Arnab, S. Jayasumana, S. Zheng, and P. H. S. Torr. Higher order conditional random fields in deep neural networks. In *European Conference on Computer Vision*, 2016. 1, 2
- [3] D. Belanger and A. McCallum. Structured prediction energy networks. In *International Conference on Machine Learning*, 2016. 1, 2
- [4] A. Blake, P. Kohli, and C. Rother. *Markov Random Fields for Vision and Image Processing*. MIT Press, 2011. 1
- [5] E. Borenstein and S. Ullman. Class-specific, top-down segmentation. In *European Conference on Computer Vision*, 2002. 5
- [6] E. Boros and P. Hammer. Pseudo-boolean optimization. *Discrete Appl. Math.*, 123:155–225, 2002. 3
- [7] L. Bottou, Y. Bengio, and Y. Le Cun. Global training of document processing systems using graph transformer networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 489–494. IEEE, 1997. 1
- [8] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001. 1
- [9] S. Chandra and I. Kokkinos. Fast, exact and multi-scale inference for semantic image segmentation with deep gaussian crfs. In *European Conference on Computer Vision*, 2016. 1, 2
- [10] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *International Conference on Learning Representations*, 2015. 1, 5
- [11] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint arXiv:1606.00915*, 2016. 1, 7
- [12] L.-C. Chen, A. Schwing, A. Yuille, and R. Urtasun. Learning deep structured models. In *International Conference on Machine Learning*, 2015. 2
- [13] Y. Chen and X. Ye. Projection onto a simplex. *arXiv preprint arXiv:1101.6081*, 2011. 4, 11
- [14] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2, 5, 7
- [15] A. Desmaison, R. Bunel, P. Kohli, P. H. S. Torr, and M. P. Kumar. Efficient continuous relaxations for dense CRF. In *European Conference on Computer Vision*, 2016. 1
- [16] J. Domke. Learning graphical model parameters with approximate marginal inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(10):2454–2467, 2013. 2
- [17] G. Ghiasi and C. C. Fowlkes. Laplacian reconstruction and refinement for semantic segmentation. In *European Conference on Computer Vision*, 2016. 1, 2, 6, 7
- [18] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 6
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 8
- [20] O. H. Jafari, O. Groth, A. Kirillov, M. Y. Yang, and C. Rother. Analyzing modular cnn architectures for joint depth prediction and semantic segmentation. In *International Conference on Robotics and Automation*, 2017. 6
- [21] V. Jampani, M. Kiefel, and P. V. Gehler. Learning sparse high dimensional filters: Image filtering, dense crfs and bilateral neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2016. 2, 4
- [22] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 2, 5
- [23] A. Kirillov, D. Schlesinger, S. Zheng, B. Savchynskyy, P. Torr, and C. Rother. Joint training of generic cnn-crf models with stochastic optimization. In *Asian Conference on Computer Vision*, 2016. 1, 2
- [24] D. Koller and N. Friedman. *Probabilistic Graphical Models*. MIT Press, 2009. 1
- [25] P. Kraehenbuehl and V. Koltun. Parameter learning and convergent inference for dense random fields. In *Proceedings of The 30th International Conference on Machine Learning*, pages 513–521, 2013. 2, 5
- [26] P. Krähenbühl and V. Koltun. Efficient inference in fully connected CRFs with gaussian edge potentials. In *Neural Information Processing Systems*. 2011. 1, 4, 5
- [27] G. Lin, C. Shen, A. van den Hengel, and I. Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2016. 1, 2, 7
- [28] Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *International Conference on Computer Vision*, 2015. 1
- [29] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 1, 2, 6
- [30] J. Peng, L. Bo, and J. Xu. Conditional neural fields. In *Advances in neural information processing systems*, pages 1419–1427, 2009. 2
- [31] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Neural Information Processing Systems*, 2015. 1
- [32] C. Rother, V. Kolmogorov, and A. Blake. "GrabCut": Interactive foreground extraction using iterated graph cuts. In *ACM Transactions on Graphics*, pages 309–314, 2004. 1
- [33] A. G. Schwing and R. Urtasun. Fully connected deep structured networks. In *arXiv preprint arXiv:1503.02351*, 2015. 1
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 1

- [35] A. Vedaldi and K. Lenc. Matconvnet – convolutional neural networks for matlab. In *Proceeding of the ACM International Conference on Multimedia*, 2015. 2, 5
- [36] V. Vineet, J. Warrell, and P. Torr. Filter-based mean-field inference for random fields with higher order terms and product label-spaces. In *European Conference on Computer Vision*, 2012. 1
- [37] P. Wang, X. Shen, Z. Lin, S. Cohen, B. Price, and A. Yuille. Towards unified depth and semantic prediction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 2, 5, 6
- [38] W. Wang, S. Fidler, and R. Urtasun. Proximal deep structured models. In *Neural Information Processing Systems*, 2016. 1, 2
- [39] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016. 7
- [40] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. In *International Conference on Computer Vision*, 2015. 1, 2, 5, 6, 7

## Supplementary Material

### Proof of Proposition 2.1

Let  $E(\mathbf{x}^*)$  and  $E(\mathbf{q}^*)$  denote the optimal values of (5) and (6), respectively. Then,

$$E(\mathbf{x}^*) = E(\mathbf{q}^*).$$

*Proof.* We will show that for any real  $\mathbf{q}$ , one can obtain a binary  $\mathbf{x}$  such that  $E(\mathbf{x}) \leq E(\mathbf{q})$ . In particular, it will be true for  $\mathbf{x}^*$  and  $\mathbf{q}^*$ , which implies  $E(\mathbf{x}^*) = E(\mathbf{q}^*)$ .

Let  $\mathbf{q}$  be given, and let  $\mathbf{x} \in \mathcal{L}^N$ . One may define  $E^m(\mathbf{x}, \mathbf{q}) = E(x_1, \dots, x_m, q_{m+1}, \dots, q_N)$  such that each  $x_i$  or  $q_i$  is a vector with entries such as  $q_{i:\lambda}$  or  $x_{i:\lambda}$ , but for each  $i$  only one value  $x_{i:\lambda}$  is non-zero (and equal to 1). Since  $E^0 = E(\mathbf{q})$  and  $E^N = E(\mathbf{x})$  it will be sufficient to find a  $\mathbf{x}$  such that  $E^m(\mathbf{x}, \mathbf{q}) \leq E^{m-1}(\mathbf{x}, \mathbf{q})$  for all  $m$ . The required  $\mathbf{x}$  will be constructed one element at a time.

The key observation is that  $E^m$  is multilinear in the  $q_i$ . Then, it follows that

$$\begin{aligned} E^{m-1}(\mathbf{x}, \mathbf{q}) &= E(x_1, \dots, x_{m-1}, q_m, \dots, q_N) \\ &= \sum_{x_m \in \mathcal{L}} q_{m:x_m} E(x_1, \dots, x_{m-1}, x_m, q_{m+1}, \dots, q_N). \end{aligned}$$

Here,  $x_m$  is treated as a variable and  $x_1, \dots, x_{m-1}$  are fixed. Since  $\sum_{x_m \in \mathcal{L}} q_{m:x_m} = 1$  there must be at least one choice of  $x_m$  such that  $E^{m-1}(\mathbf{x}, \mathbf{q}) \geq E(x_1, \dots, x_m, q_{m+1}, \dots, q_m) = E^m(\mathbf{x}, \mathbf{q})$ .  $\square$

### Simplex Projection

As mentioned in the paper we want to project our values onto the simplex  $\Delta^L$  satisfying  $\sum_{\lambda \in \mathcal{L}} q_{i:\lambda} = 1$  and  $0 \leq q_{i:\lambda} \leq 1$  for each step in our gradient descent algorithm. This is done following the method presented by Chen *et al.* [13] which is summarized in Algorithm 2. Note that this projection is done individually for each pixel  $i$ , for derivation of this algorithm we refer to the article.

1. Sort  $\tilde{q}_i \in \mathbb{R}^L$  in ascending order and set  $k = L - 1$
2. Compute  $t_i = \frac{\sum_{j=k+1}^L \tilde{q}_{i:j} - 1}{L - k}$ , If  $t_i \geq \tilde{q}_{i:k}$  set  $\hat{t} = t_i$  and go to step 4, otherwise set  $k \leftarrow k - 1$  and redo step 2, if  $k = 0$  go to step 3
3. Set  $\hat{t} = \frac{\sum_{j=1}^L \tilde{q}_{i:j} - 1}{L}$
4. Return  $\mathbf{q}_i$ , where  $q_{i:\lambda} = \max(\tilde{q}_{i:\lambda} - \hat{t}, 0)$ ,  $\lambda \in \mathcal{L}$

Algorithm 2. Projection of  $\tilde{q}_i \in \mathbb{R}^L$  onto the simplex  $\Delta^L$  satisfying  $\sum_{\lambda \in \mathcal{L}} q_{i:\lambda} = 1$  and  $0 \leq q_{i:\lambda} \leq 1$ .

In reality we use a leaky version of the last step of the projection algorithm, *i.e.* instead of the  $\max(\cdot, 0)$  operator

we use the following function

$$f_\alpha(\tilde{q}_{i:\lambda}) = \begin{cases} \tilde{q}_{i:\lambda} - \hat{t} & 0 \leq \tilde{q}_{i:\lambda} - \hat{t} \\ \alpha (\tilde{q}_{i:\lambda} - \hat{t}) & \tilde{q}_{i:\lambda} - \hat{t} < 0 \end{cases} \quad (12)$$

where  $\alpha$  is a parameter controlling the level of leakage. Note that for  $\alpha = 0$  we get the strict  $\max(\cdot, 0)$  operator. As previously mentioned, the projection is done individually for each pixel. It can be described as a function  $\mathbf{f}(\tilde{\mathbf{q}}) : \mathbb{R}^L \rightarrow \mathbb{R}^L$ , which Jacobian has the elements

$$\frac{\partial f_\lambda}{\partial \tilde{q}_\mu} = \begin{cases} \begin{cases} \alpha \left(1 - \frac{\partial \hat{t}}{\partial \tilde{q}_\mu}\right) & \tilde{q}_\mu - \hat{t} < 0 \\ 1 - \frac{\partial \hat{t}}{\partial \tilde{q}_\mu} & \tilde{q}_\mu - \hat{t} \geq 0 \end{cases} & \mu = \lambda \\ \begin{cases} -\alpha \frac{\partial \hat{t}}{\partial \tilde{q}_\mu} & \tilde{q}_\mu - \hat{t} < 0 \\ -\frac{\partial \hat{t}}{\partial \tilde{q}_\mu} & \tilde{q}_\mu - \hat{t} \geq 0 \end{cases} & \mu \neq \lambda \end{cases} \quad (13)$$

where  $\frac{\partial \hat{t}}{\partial \tilde{q}_\mu} = \frac{1}{L-k}$  if  $\tilde{q}_\mu > \hat{t}$  and 0 otherwise ( $L$  and  $k$  defined as in Algorithm 2). Knowing the Jacobian, the error derivatives with respect to the input can be computed during back propagation. The reason for introducing the leaky version is to avoid error derivatives becoming zero during back propagation. A non-zero  $\alpha$  was found to facilitate the training process considerably with the drawback that we do not necessarily satisfy the constraints in (6). However, one could set  $\alpha$  to zero during inference to strictly satisfy (6).

### Convergence analysis

Since the forward operation of the CRF-Grad layer performs gradient descent, we are interested in knowing how many iterations are needed to converge. In Fig. 9 we have plotted the CRF energy as a function of the number of iterations on the WEIZMANN HORSE data set for the different models. In addition Fig. 10 shows the mean IoU as a function of the number of iterations for the WEIZMANN HORSE test data. Both of these results point to the fact that it is sufficient to run the model for about five iterations and that running it further would not increase the result considerably.

It also might be interesting to investigate the intermediate states  $\mathbf{q}^t$ . In Fig. 11 the intermediate states of the CRF-Grad layers is shown. This figure gives a good indication on affect the CRF-Grad layer has. As can be seen in the figure, each step refines the segmentation slightly, removing spurious outlier pixels classified as horse, in addition to refining the boundaries slightly.

Also included in the supplementary material is short movie, crfgd2\_mov.avi, showing how the CRF energy decreases with each step of our projected gradient descent algorithm. For this example a Potts model is used and the projected gradient descent solution is compared to the globally optimal solution obtained by graph cut.

## Error Derivatives for the CRF-Grad layer

In this section, we will explicitly formulate the error derivative necessary to train our deep structure model jointly. The notation used in the section is not very strict. Derivatives, gradients and jacobians are all referred to as derivatives. Denoting the output of our CRF-Grad layer  $\mathbf{y}$  we need expressions for the derivatives  $\frac{\partial \mathbf{y}}{\partial \mathbf{z}}$ , where  $\mathbf{z}$  is the output from the CNN and hence also the input to the CRF-Grad layer. In addition we need to calculate  $\frac{\partial \mathbf{y}}{\partial \mathbf{w}_u}$ ,  $\frac{\partial \mathbf{y}}{\partial \mathbf{w}_s}$  and  $\frac{\partial \mathbf{y}}{\partial \mathbf{w}_b}$  to be able to update the weights of the layer. To simplify the notation we abbreviate the update step by  $\mathbf{q}^{t+1} = f(\mathbf{q}^t, \mathbf{z}, I, \mathbf{w})$ . Note that the output  $\mathbf{y} = \mathbf{q}^T$  where  $T$  is the total number of iterations for the RNN. We have

$$\frac{\partial \mathbf{y}}{\partial \mathbf{w}_u} = \frac{\partial \mathbf{y}}{\partial \mathbf{q}^T} \frac{\partial f(\mathbf{q}^{T-1})}{\partial \mathbf{w}_u} + \dots + \frac{\partial \mathbf{y}}{\partial \mathbf{q}^1} \frac{\partial f(\mathbf{q}^0)}{\partial \mathbf{w}_u} \quad (14)$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{w}_s} = \frac{\partial \mathbf{y}}{\partial \mathbf{q}^T} \frac{\partial f(\mathbf{q}^{T-1})}{\partial \mathbf{w}_s} + \dots + \frac{\partial \mathbf{y}}{\partial \mathbf{q}^1} \frac{\partial f(\mathbf{q}^0)}{\partial \mathbf{w}_s} \quad (15)$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{w}_b} = \frac{\partial \mathbf{y}}{\partial \mathbf{q}^T} \frac{\partial f(\mathbf{q}^{T-1})}{\partial \mathbf{w}_b} + \dots + \frac{\partial \mathbf{y}}{\partial \mathbf{q}^1} \frac{\partial f(\mathbf{q}^0)}{\partial \mathbf{w}_b} \quad (16)$$

$$\frac{\partial \mathbf{y}}{\partial \mathbf{z}} = \frac{\partial \mathbf{y}}{\partial \mathbf{q}^0} \frac{\partial \mathbf{q}^0}{\partial \mathbf{z}} + \frac{\partial \mathbf{y}}{\partial \psi_u} \frac{\partial \psi_u}{\partial \mathbf{z}}, \quad (17)$$

where  $\psi_u$  denote the unary part of the CRF energy function. Note that

$$\frac{\partial \mathbf{y}}{\partial \mathbf{q}^{t-1}} = \frac{\partial \mathbf{y}}{\partial \mathbf{q}^t} \frac{\partial f(\mathbf{q}^{t-1})}{\partial \mathbf{q}^{t-1}} \quad (18)$$

To be able to calculate these we need the derivatives of the function  $f$  with respect to  $\mathbf{q}^t$ ,  $\mathbf{w}_u$ ,  $\mathbf{w}_s$  and  $\mathbf{w}_b$ . We denote the spatial and bilateral filtering operations as  $\psi_s * \mathbf{q}^t$  and  $\psi_b * \mathbf{q}^t$  respectively. An update step can then be written as

$$\begin{aligned} \mathbf{q}^{t+1} &= \text{Proj}_{\Delta^L}(\mathbf{q}^t - \gamma(\psi_u + \psi_s * \mathbf{q}^t + \psi_b * \mathbf{q}^t)) \\ &= \text{Proj}_{\Delta^L}(\tilde{\mathbf{q}}^{t+1}). \end{aligned}$$

And the aforementioned derivatives become

$$\begin{aligned} \frac{\partial f}{\partial \mathbf{q}^t} &= \text{Proj}'_{\Delta^L}(\tilde{\mathbf{q}}^{t+1}) \\ &\cdot (1 - \gamma(\psi_u + \psi_s * \mathbf{q}^t + \psi_b * \mathbf{q}^t)) \\ &\cdot \left( \frac{\partial(\psi_s * \mathbf{q}^t)}{\partial \mathbf{q}^t} + \frac{\partial(\psi_b * \mathbf{q}^t)}{\partial \mathbf{q}^t} \right), \end{aligned} \quad (19)$$

for  $\mathbf{q}^t$ , and for the weights

$$\begin{aligned} \frac{\partial f}{\partial \mathbf{w}_u} &= \text{Proj}'_{\Delta^L}(\tilde{\mathbf{q}}^{t+1}) \cdot \\ &(-\gamma(\psi_u + \psi_s * \mathbf{q}^t + \psi_b * \mathbf{q}^t)) \cdot \frac{\partial \psi_u}{\partial \mathbf{w}_u}, \end{aligned} \quad (20)$$

$$\begin{aligned} \frac{\partial f}{\partial \mathbf{w}_s} &= \text{Proj}'_{\Delta^L}(\tilde{\mathbf{q}}^{t+1}) \cdot \\ &(-\gamma(\psi_u + \psi_s * \mathbf{q}^t + \psi_b * \mathbf{q}^t)) \cdot \frac{\partial(\psi_s * \mathbf{q}^t)}{\partial \mathbf{w}_s}, \end{aligned} \quad (21)$$

$$\begin{aligned} \frac{\partial f}{\partial \mathbf{w}_b} &= \text{Proj}'_{\Delta^L}(\tilde{\mathbf{q}}^{t+1}) \cdot \\ &(-\gamma(\psi_u + \psi_s * \mathbf{q}^t + \psi_b * \mathbf{q}^t)) \cdot \frac{\partial(\psi_b * \mathbf{q}^t)}{\partial \mathbf{w}_b}, \end{aligned} \quad (22)$$

Note that  $\frac{\partial(\psi_s * \mathbf{q}^t)}{\partial \mathbf{q}^t}$ ,  $\frac{\partial(\psi_b * \mathbf{q}^t)}{\partial \mathbf{q}^t}$ ,  $\frac{\partial(\psi_s * \mathbf{q}^t)}{\partial \mathbf{w}_s}$  and  $\frac{\partial(\psi_b * \mathbf{q}^t)}{\partial \mathbf{w}_b}$  can be calculated using the backward routines for a standard convolutional layer and bilateral filtering layer described in the main paper.

## Additional Results

In this section we present some additional results from the different datasets. Class-wise results for the CITYSCAPES dataset is presented in Table 4. For the WEIZMANN HORSE dataset these results can be seen in Fig. 12, for the CITYSCAPES dataset in Fig. 13 and for the NYU V2 dataset in Fig. 14.



	road	sidewalk	building	wall	fence	pole	traffic light	traffic sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motorcycle	bicycle
LRR-4x	<b>97.9</b>	<b>81.5</b>	<b>91.4</b>	50.5	52.7	59.4	66.8	72.7	92.5	<b>70.1</b>	<b>95</b>	81.3	<b>60.1</b>	<b>94.3</b>	<b>51.2</b>	67.7	54.6	55.6	69.6
CRF-Grad	<b>97.9</b>	81.3	<b>91.4</b>	<b>51.4</b>	<b>52.9</b>	<b>60.65</b>	<b>68.1</b>	<b>73.5</b>	<b>92.6</b>	69.1	<b>95</b>	<b>81.6</b>	59.5	94.1	45.5	<b>68.4</b>	<b>56</b>	<b>56</b>	<b>70.2</b>

Table 4. Class-wise results compairing our model to the baseline, LRR. The results are presented as IoU (%).

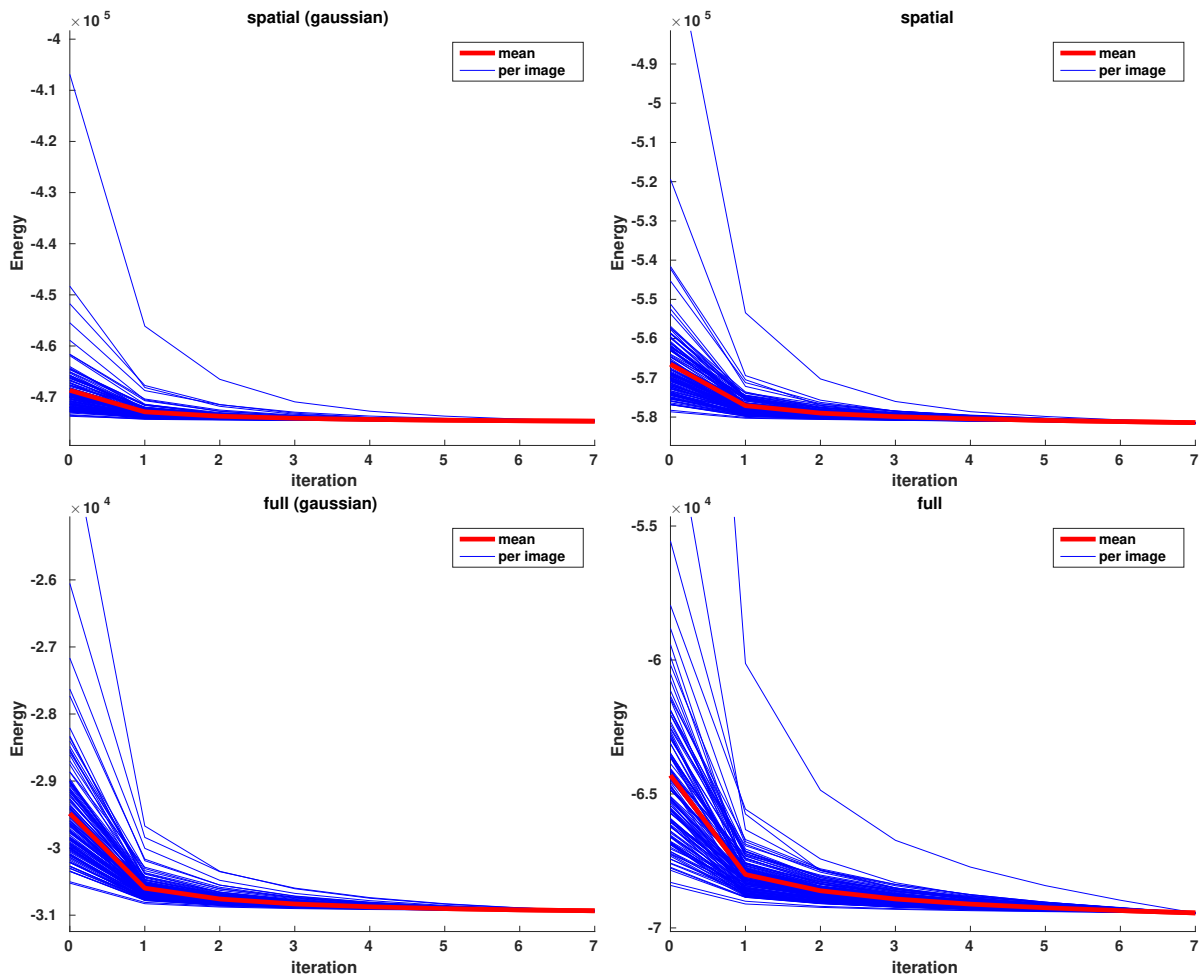


Figure 9. CRF energy, as defined in (6), plotted as a function of iterations of the gradient descent method. These energies are calculated while running the trained models on the WEIZMANN HORSE dataset. The thin blue lines show the imagewise energy change while the thicker red line shows the mean. For these calculations the leak factor was set to zero, meaning that the solutions satisfy the constraints of (6). Note that, for presentation purposes, all energies have been normalized to have the same final energy.

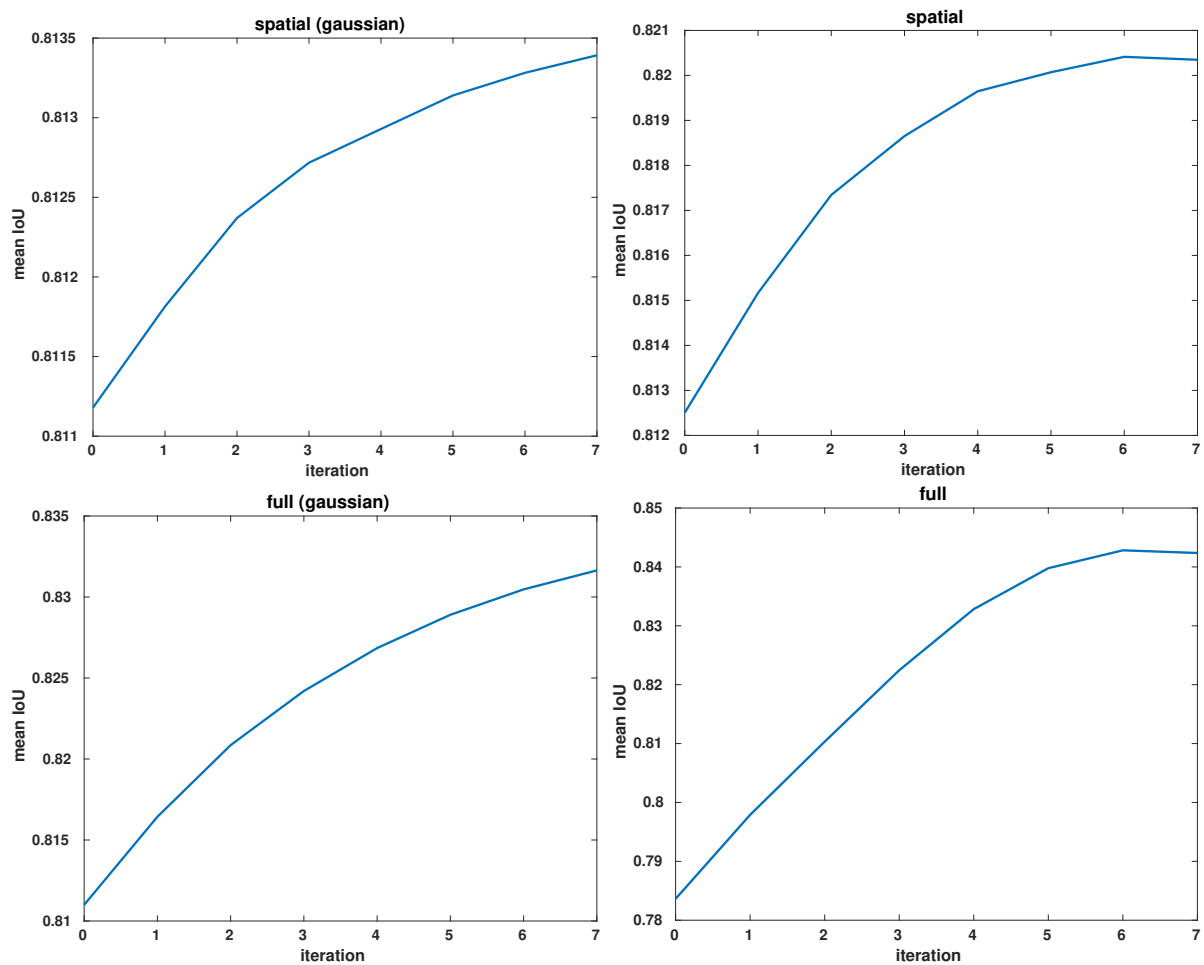


Figure 10. Mean IoU plotted as a function of iterations for the WEIZMANN HORSE test data. Note that, for the non-gaussian versions of the model, the performance of the layer does not increase after six iterations. The number of iterations were set to five during training.




Figure 11. Visualization of intermediate states of the CRF-Grad layer for the WEIZMANN HORSE dataset. From left: Original image,  $z_{\text{horse}}$  (CNN output),  $q_{\text{horse}}^1$ ,  $q_{\text{horse}}^2$ ,  $q_{\text{horse}}^3$ ,  $q_{\text{horse}}^4$ ,  $y_{\text{horse}}$  (CRF-Grad output) and finally ground truth. Note how each step of the projected gradient descent algorithm refines the segmentation slightly, removing spurious outlier pixels classified as horse.




Figure 12. Qualitative results on the WEIZMANN HORSE dataset. From left: Original image, Ground truth, FCN-8s segmentation results, CRF-Grad segmentation results (spatial kernel only), CRF-Grad segmentation results (spatial and bilateral kernels). Note how adding the CRF-Grad layer gives more refined segmentations, removing spurious outlier pixels previously classified as horse.




Figure 13. Qualitative results on the CITYSCAPES dataset. From left: Original image, Ground truth, LRR segmentation results, LRR + CRF-Grad segmentation results. Note that the poles are better segmented with the additional CRF-Grad layer.





Figure 14. Qualitative results on the NYU V2 dataset. From left: Original image, Ground truth, FCN-8s segmentation results, FCN-8s + CRFasRNN segmentation results, FCN-8s + CRF-Grad segmentation results. Note that the CRF-Grad captures the shape of the object instances better compared to the baselines. This effect is perhaps most pronounced for the paintings hanging on the walls.